

Deep Linking Metrics and Traces

with OpenTelemetry, OpenMetrics,
Prometheus and M3

San Diego, 2019-11-21
Rob Skillington

Who



Rob Skillington


CTO at Chronosphere

Previously M3 and M3DB technical lead at Uber

OpenMetrics Contributor

Let's talk about

1. State of Monitoring: Logs, Metrics, Traces
2. Combining Logs, Metrics and Traces Today
3. Deep Linking Metrics and Traces



1 State of Monitoring: Logs, Metrics, Traces

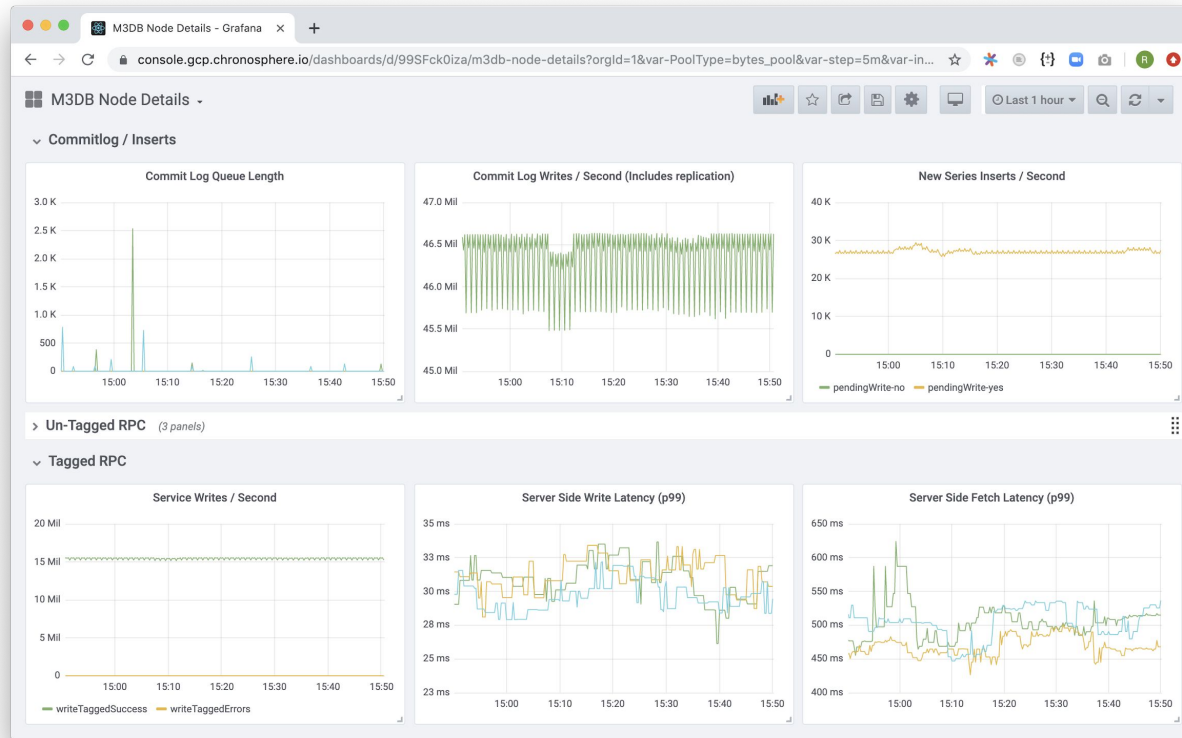


Logs

The screenshot shows a web browser window with multiple tabs. The active tab is titled "[eCommerce] Revenue Dashboard". The browser's address bar shows the URL: `demo.elastic.co/app/kibana#/dashboard/722b74f0-b882-11e8-a6d9-e546fe2bba5f?_g=(refreshInterval:(pause:f,value:900000),time:(from:now-7d,to:now))&a=(description:'Analyze%20...'`. The dashboard header includes "Full screen" and "Share" buttons, a search bar with "# Search", and a filter dropdown set to "Lucene". The time range is "Last 7 days" with "Show dates" and "Refresh" buttons. Below the header, the section is titled "[eCommerce] Orders" and shows a table with 15 rows of order data. The table has columns for "Time", "category", "sku", "taxful_total_price", and "total_quantity". The "Time" column includes a dropdown arrow. The "category" column contains hierarchical category names. The "sku" column lists multiple SKUs for each order. The "taxful_total_price" column includes a magnifying glass icon. The "total_quantity" column shows the number of items per order. The table is paginated, showing "1-50 of 1046" items.


Time	category	sku	taxful_total_price	total_quantity
> Nov 18, 2019 @ 20:55:12.000	Women's Clothing, Women's Shoes	Z00707507075, Z00246402464, Z00226802268, Z00343503435	\$139.96	4
> Nov 18, 2019 @ 20:50:53.000	Men's Clothing	Z00473704737, Z00121501215	\$45.98	2
> Nov 18, 2019 @ 20:35:02.000	Women's Shoes, Women's Clothing	Z00673606736, Z00161801618	\$88.98	2
> Nov 18, 2019 @ 20:04:48.000	Women's Shoes, Women's Accessories	Z00242702427, Z00090000900	\$70.98	2
> Nov 18, 2019 @ 19:59:02.000	Men's Clothing	Z00589505895, Z00575405754	\$42.98	2
> Nov 18, 2019 @ 19:51:50.000	Women's Clothing, Women's Shoes	Z00490204902, Z00025000250	\$45.98	2
> Nov 18, 2019 @ 19:50:24.000	Men's Shoes, Men's Clothing	Z00400004000, Z00519305193, Z00482004820, Z00540305403	\$300.96	4
> Nov 18, 2019 @ 19:50:24.000	Men's Clothing	Z00419604196, Z00559705597	\$39.98	2
> Nov 18, 2019 @ 19:24:29.000	Men's Shoes, Men's Accessories	Z00520305203, Z00462204622	\$66.96	2
> Nov 18, 2019 @ 19:17:17.000	Women's Shoes, Women's Clothing	Z00216502165, Z00327503275	\$78.98	2
> Nov 18, 2019 @ 19:14:24.000	Men's Shoes, Men's Clothing	Z00257002570, Z00455404554	\$85.98	2
> Nov 18, 2019 @ 19:08:38.000	Men's Clothing	Z00547905479, Z00583305833	\$32.98	2
> Nov 18, 2019 @ 18:55:41.000	Women's Clothing	Z00341103411, Z00648406484	\$60.98	2
> Nov 18, 2019 @ 18:52:48.000	Women's Clothing	Z00100901009, Z00235102351	\$53.98	2
> Nov 18, 2019 @ 18:51:22.000	Men's Clothing	Z00575305753, Z00540605406	\$58.98	2
> Nov 18, 2019 @ 18:39:50.000	Women's Clothing, Women's Shoes	Z00266902669, Z00244202442	\$105.98	2
> Nov 18, 2019 @ 18:31:12.000	Men's Clothing	Z00279702797, Z00573705737	\$48.98	2

Metrics



Metrics

Google trends of popular metrics formats



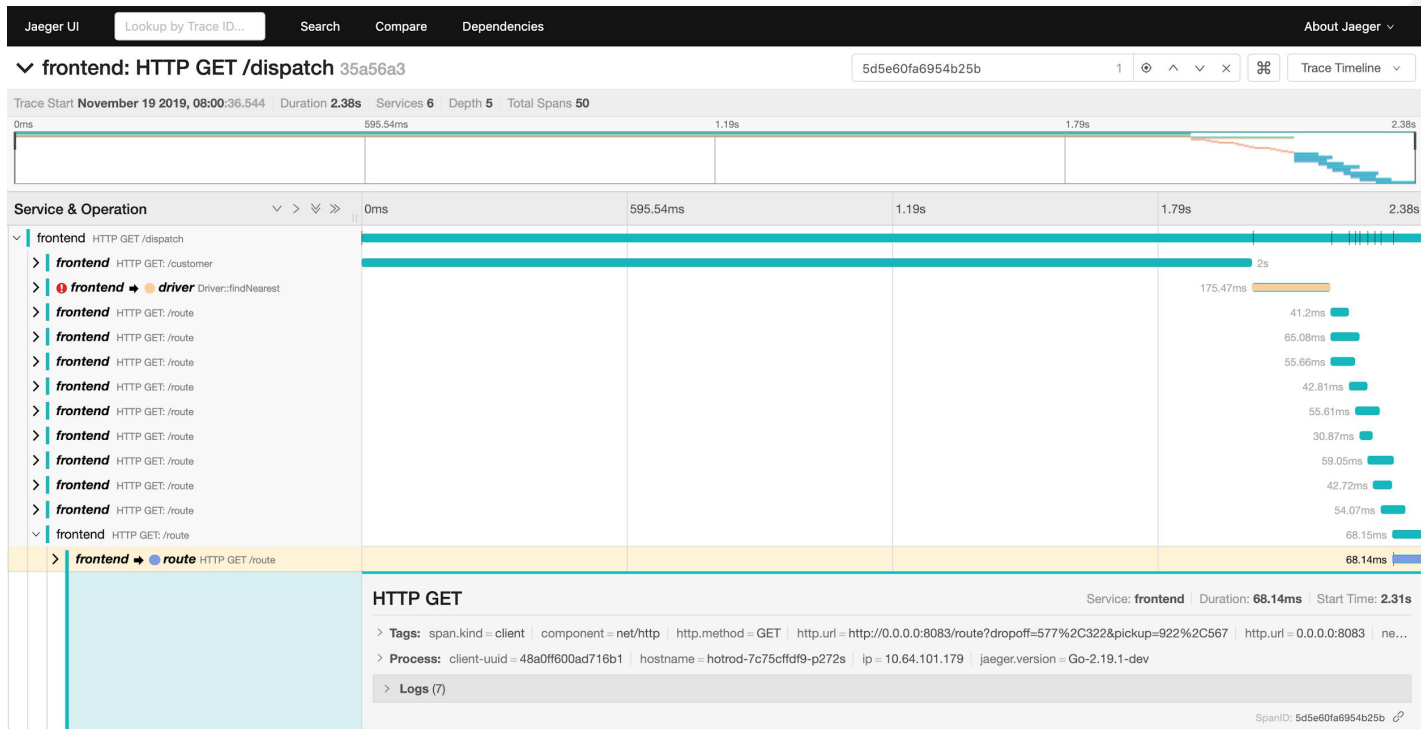
Metrics

Numbers shared at PromCon 2019

2016	2017	2018	2019
~2,800	~16,000	~54,000	~242,000




Tracing



Tracing


Star history




A perspective on... Logs



A perspective on... Traces



A perspective on... Metrics

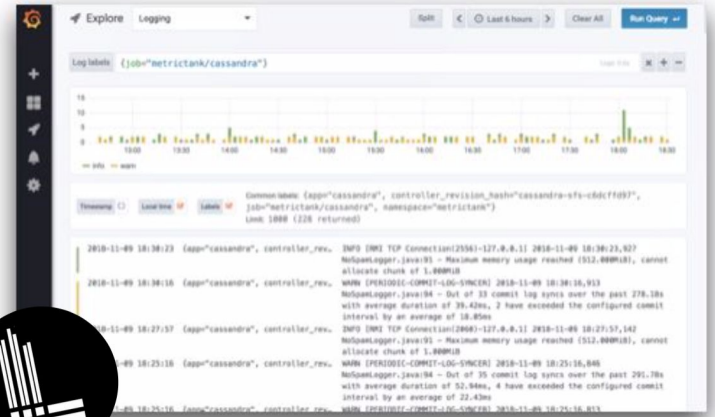
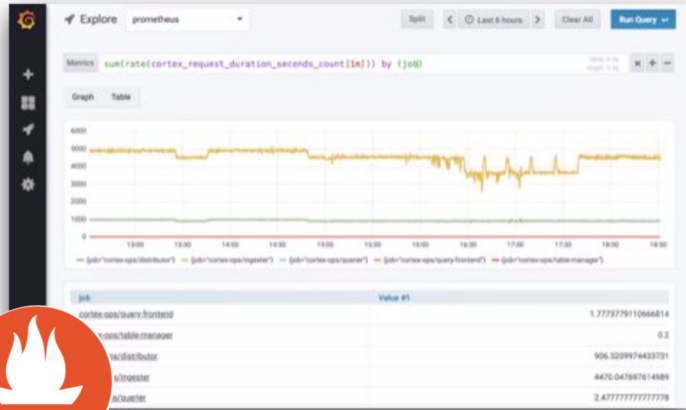


2 Combining Logs, Metrics, and Traces Today



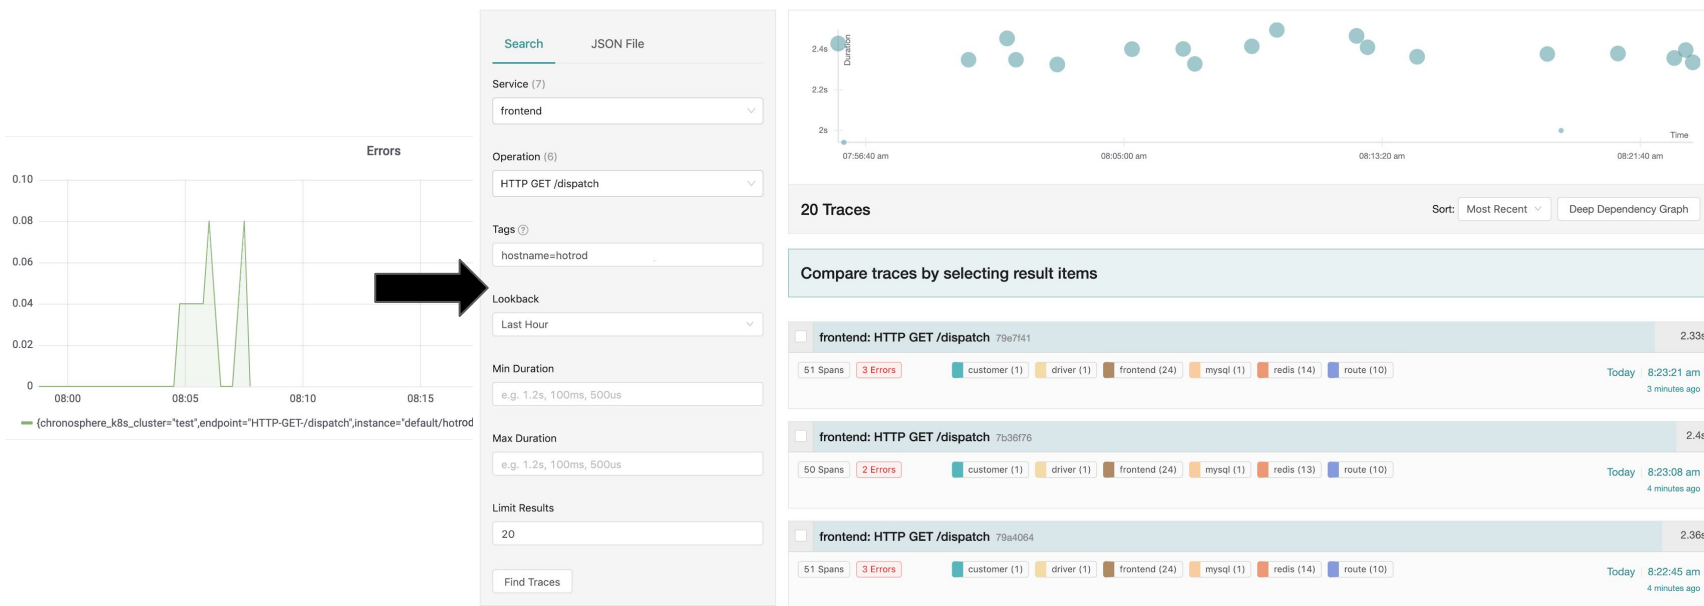
Current integrations

Increasingly, more Observability platforms providing two or more signals (logs, metrics, traces)



Integrations - Taking a closer look

If you take a closer look at how jumping between metrics and traces is today, they are generally linked by common set of labels and time window.



Integrations - Taking a closer look

This narrows down search space in terms of time window and labels, but:

- Querying metrics with `sum(...)` or any other aggregation will drop tags
 - Context lost for jumping from metrics to traces
- Only “magical” when you store every trace
 - For a lot of users is prohibitively expensive.
- When sampling, the chances of having the right trace is low
 - Especially debugging edge cases - P99, or one error in a thousand.

Wouldn't it be nice if?

Go straight from the metric datapoint to one of the traces for a request that comprised that **exact datapoint**.

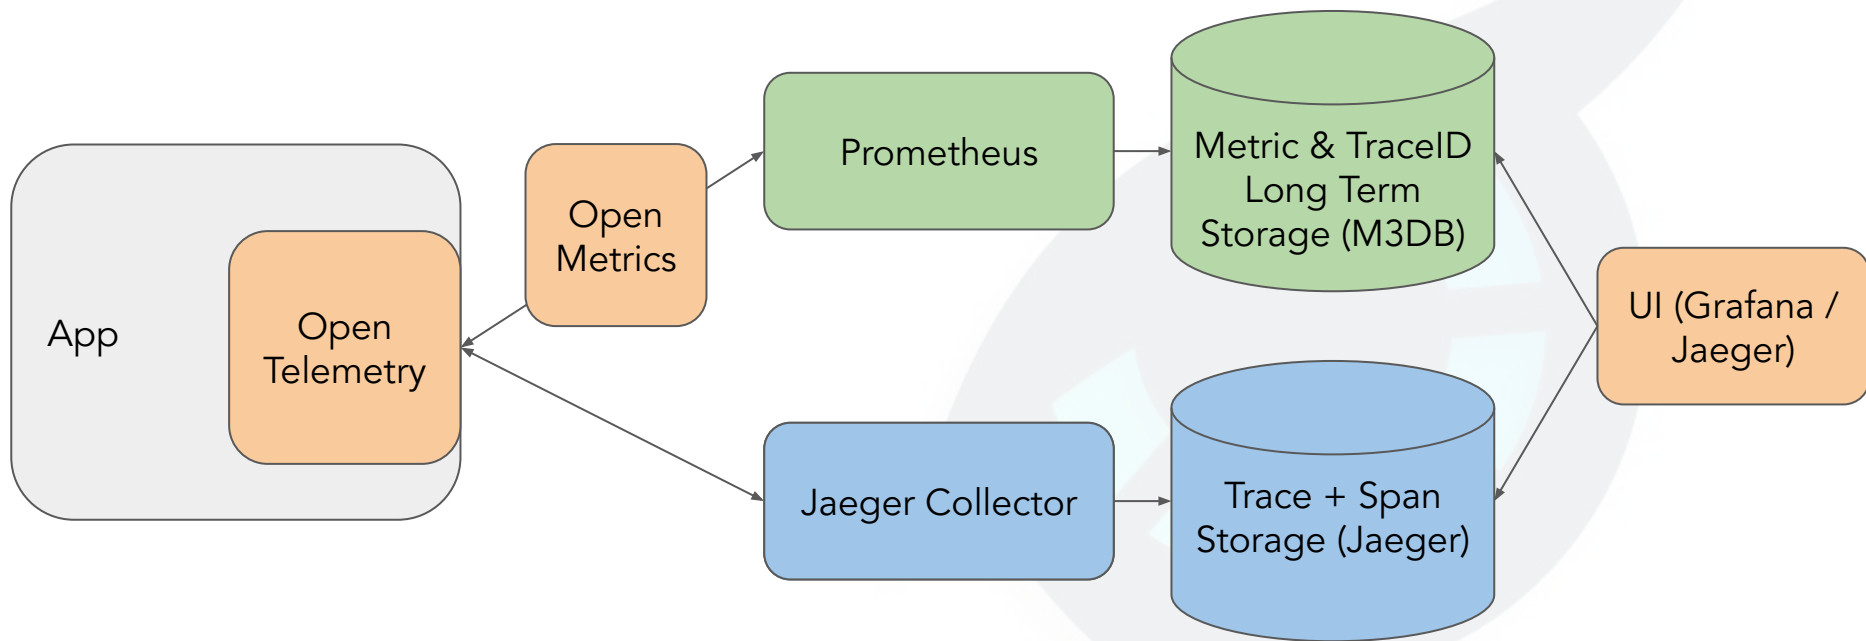


3 Deep Linking Metrics and Traces



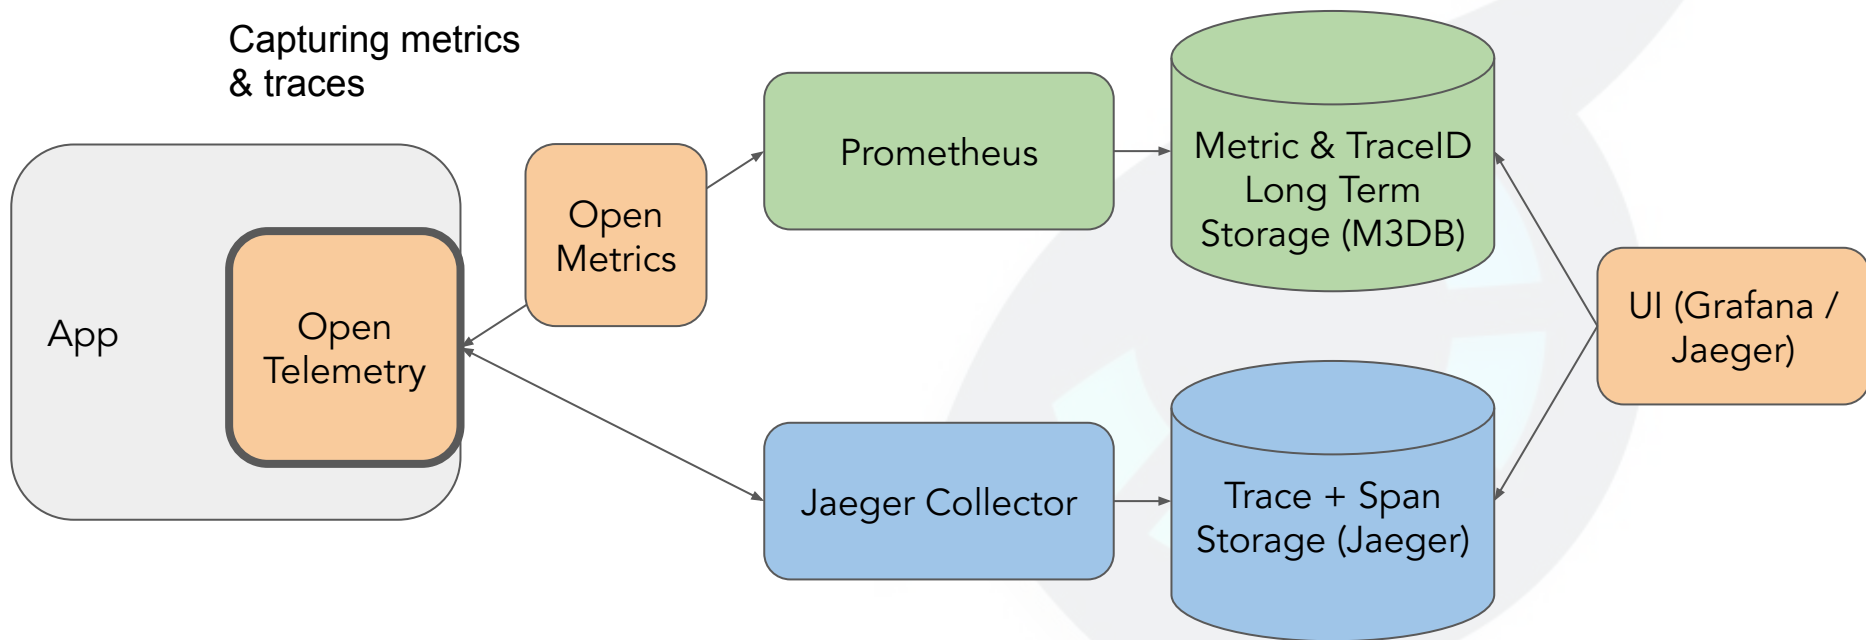
Demo

Putting it altogether



What is OpenTelemetry?

Capturing metrics
& traces



OpenTelemetry: Instrumentation SDK

```
jobsQueuedGauge := meter.NewFloat64Gauge("jobs_queued",  
    metric.WithDescription("The number of jobs currently queued"))  
  
err := tracer.WithSpan(ctx, "jobEnqueue", func(ctx context.Context) error {  
    jobsTotal, err := jobQueue.Enqueue(job)  
    if err != nil {  
        return err  
    }  
    jobsQueuedGauge.Set(ctx, jobsTotal)  
})
```



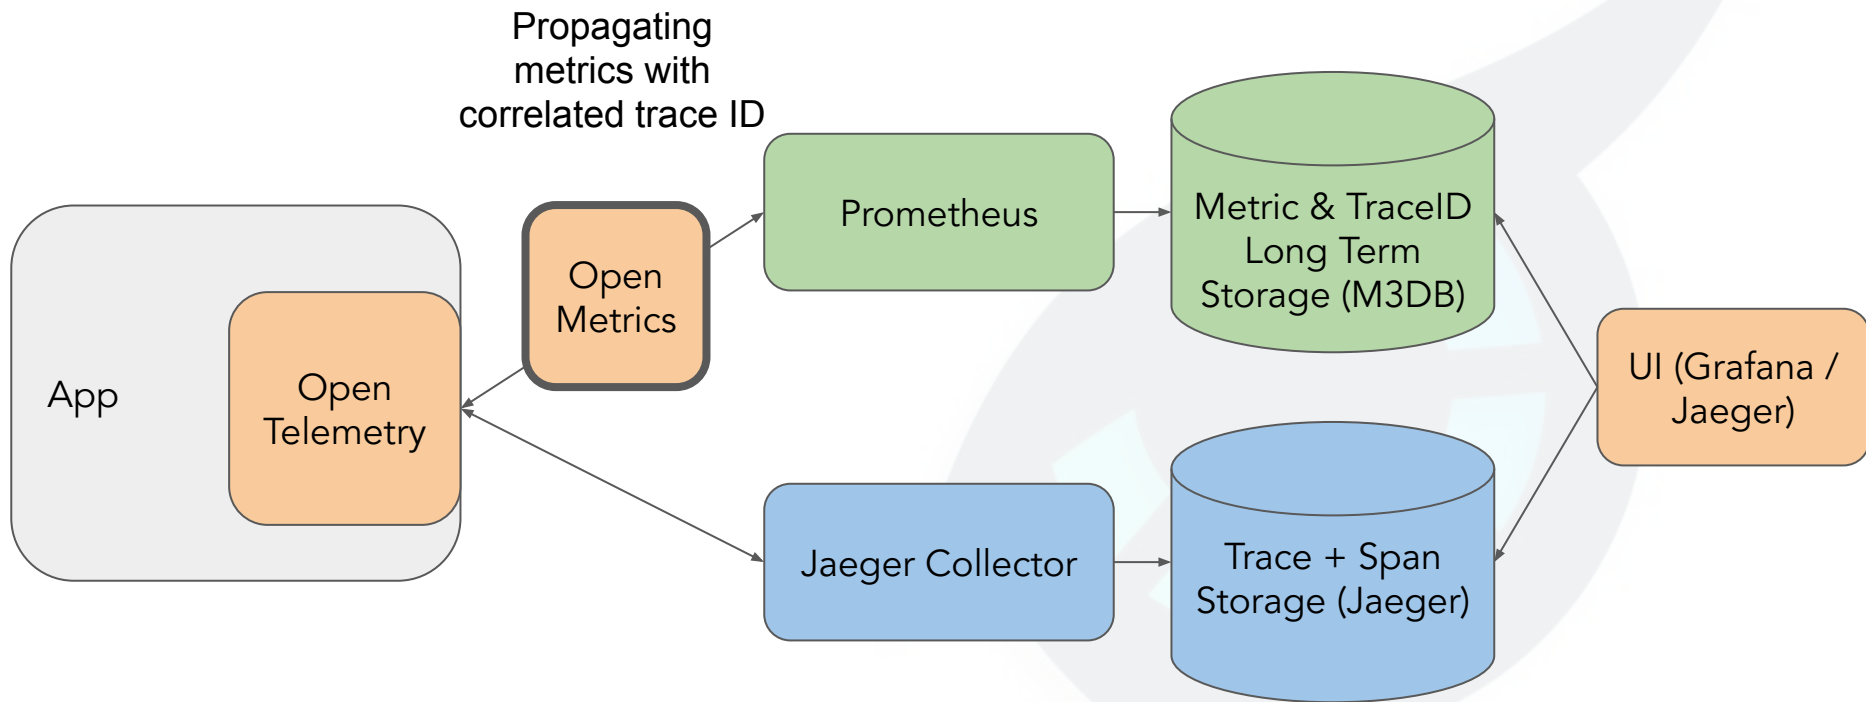
OpenTelemetry: Instrumentation SDK

```
jobsQueuedGauge := meter.NewFloat64Gauge("jobs_queued",
    metric.WithDescription("The number of jobs currently queued"))

err := tracer.WithSpan(ctx, "jobEnqueue", func(ctx context.Context) error {
    jobsTotal, err := jobQueue.Enqueue(job)
    if err != nil {
        return err
    }
    jobsQueuedGauge.Set(ctx, jobsTotal)
})
```



What is OpenMetrics?

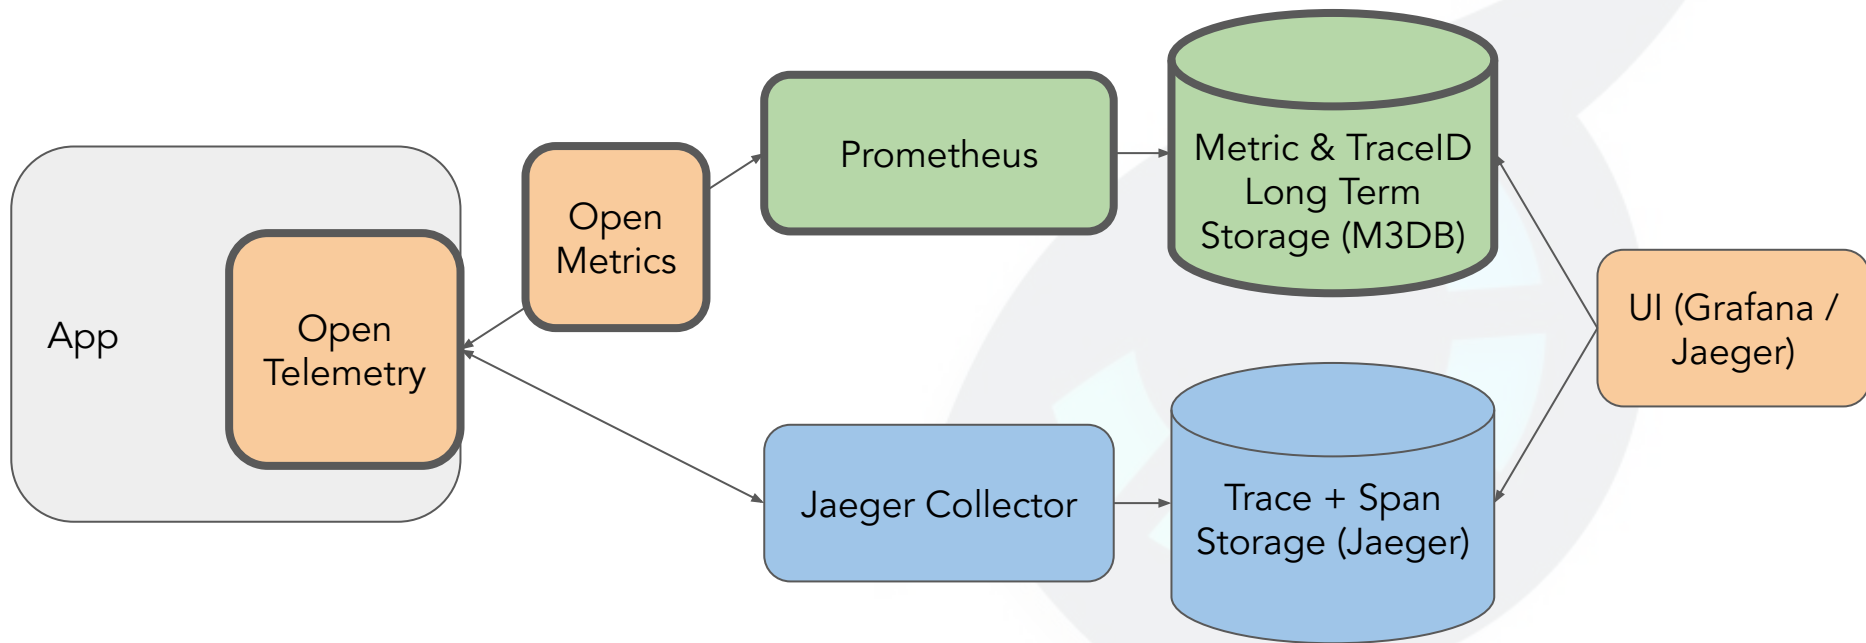


OpenMetrics: Extended Prometheus exposition

```
# HELP http_requests_total http_requests
# TYPE http_requests_total counter
http_requests_total{endpoint="/search",status_code="2xx"} 1725 # {trace_id="b096e71d..."} 1
http_requests_total{endpoint="/search",status_code="4xx"} 4 # {trace_id="944a6d97..."} 1
http_requests_total{endpoint="/search",status_code="5xx"} 27 # {trace_id="50785260..."} 1
http_request_latency_bucket{endpoint="/search",le="0.1"} 7 # {trace_id="7f78deda..."} 1
http_request_latency_bucket{endpoint="/search",le="0.2"} 7 # {trace_id="5ad53ac9..."} 1
http_request_latency_bucket{endpoint="/search",le="0.3"} 7 # {trace_id="c78493ec..."} 1
...
```



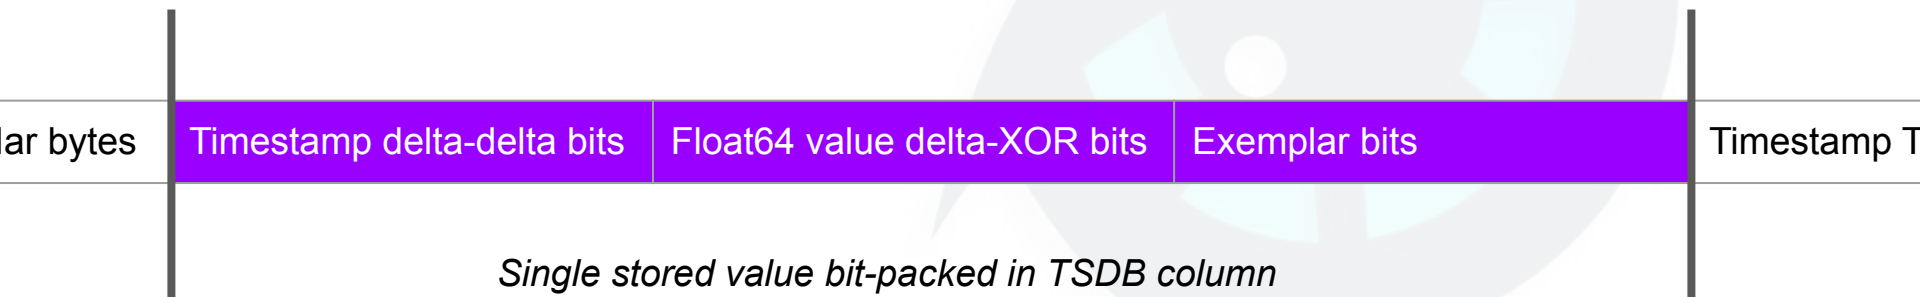
Putting it all together



Prometheus and M3

Prometheus scrapes the exemplar, keeps it locally in memory, then remote writes it to M3.

M3DB stores the trace ID next to the metric timestamp and float value.

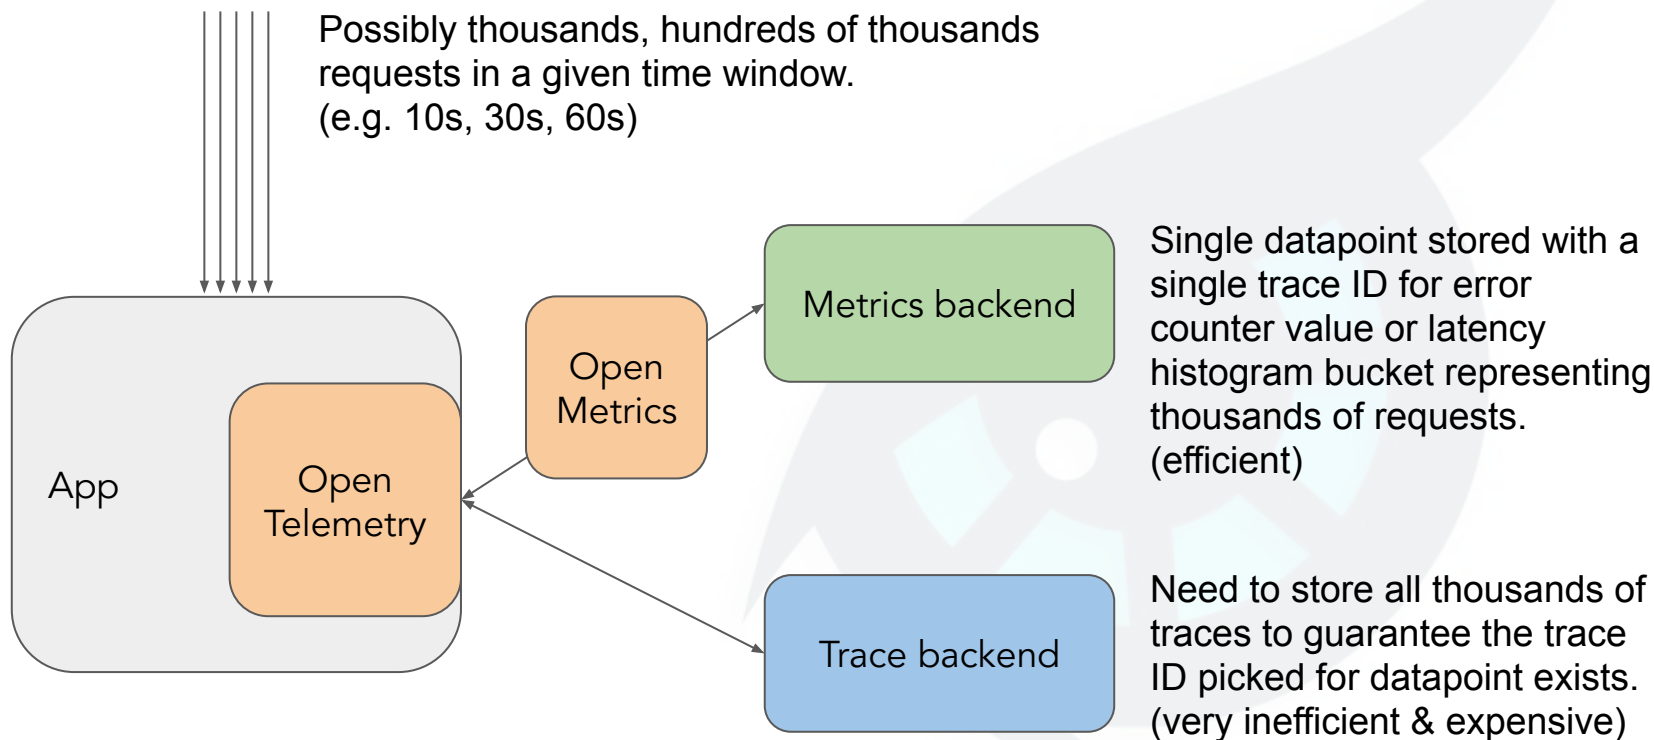


How do we query it?

When querying the data M3 query makes sure to keep at least one representative exemplar per datapoint as part of the result (even after applying `sum(...)`, `histogram_quantile(...)`, etc)

```
{
  "metric": {
    "chronosphere_k8s_cluster": "test",
    "endpoint": "HTTP-GET-/customer",
    "instance": "default/hotrod-7c75cffdf9-p272s",
    "job": "hotrod",
    "status_code": "5xx"
  },
  "values": [
    [
      1574178960,
      "5",
      "trace_id:ac51ab1117abfdc:d4c6a84f5833856:3638a233f5a36b71:1"
    ]
  ]
}
```

Efficiency/Scalability?



Wouldn't it be nice if?

200 Status Code

{trace: 024253eb-6be0-...}
{trace: 841e6da2-8694-...}
{trace: f7e33019-abc8-...}
{trace: 7b2a9954-e213-...}
{trace: d37ce450-a463-...}
{trace: b78fe85b-a508-...}
{trace: b10964a4-a4af-...}
{trace: 1cdfcb7a-1849-...}
{trace: 3bb247b2-89ec-...}
{trace: 2bce5524-905e-...}

400 Status Code

{trace: a0eb52dc-8a3e-...}
{trace: f86aa034-b7c5-...}
{trace: 6aa9d08f-6632-...}
{trace: 1be7ef05-9985-...}

500 Status Code

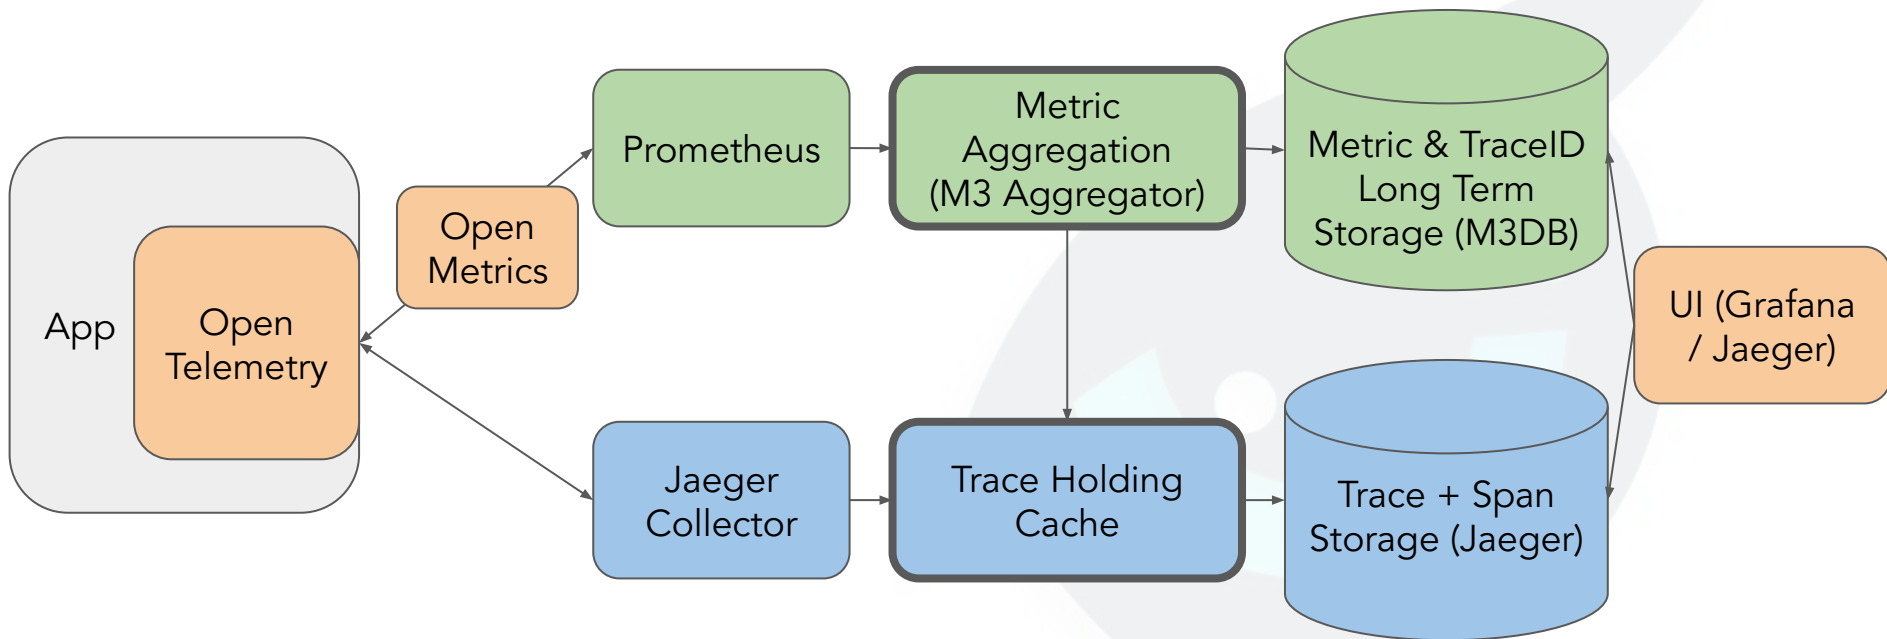
{trace: a22476ff-b177-...}

Guaranteeing one representative trace stored?

Using metric aggregation to determine sampled traces has many upsides:

- Aggregate across time to collect traces at **useful time intervals**.
- Metric tags are great at capturing all **unique combinations**. eg:
Error/Success, ErrorStatusCode, LatencyBucket.
 - This also ensures unique combinations of traces.
- Allows maintaining **direct link** between metric datapoint and trace ID.

Putting it altogether (again)



Where can I get this? (hint: upstream in progress)

Current end-to-end demo at:

<https://github.com/chronosphereio/demo-deeplink-metrics-traces>

Merged: Add exemplar support to OpenMetrics:

<https://github.com/prometheus/prometheus/pull/6292>

Open(needs discussion): Store exemplars in Prometheus memory, forward on remote write:

<https://github.com/prometheus/prometheus/pull/6309>

Open(helping review): OpenMetrics/Prometheus exporter PR for OpenTelemetry:

<https://github.com/open-telemetry/opentelemetry-go/pull/334>

Where can I get this?

OpenMetrics <https://github.com/OpenObservability/OpenMetrics>

OpenTelemetry <https://github.com/open-telemetry/opentelemetry-specification>

Prometheus <https://github.com/prometheus/prometheus>

M3 <https://github.com/m3db/m3>

Grafana <https://github.com/grafana/grafana>

Talk demo <https://github.com/chronosphereio/demo-deeplink-metrics-traces>

Thank you and Q&A

Come say hi! Booth SE62.



chronosphere



<http://bit.ly/m3slack>